arXiv:1701.01830v2 [math.OC] 18 Jan 2017

Parallel algorithms and probability of large deviation for stochastic optimization problems

Pavel Dvurechensky [1], *Alexander Gasnikov* [2], *Anastasia Lagunovskaya* [3]

Abstract

We consider convex stochastic optimization problems under different assumptions on the properties of available stochastic subgradient. It is known that, if the value of the objective function is available, one can obtain, in parallel, several independent approximate solutions in terms of the objective residual expectation. Then, choosing the solution with the minimum function value, one can control the probability of large deviation of the objective residual. On the contrary, in this short paper, we address the situation, when the value of the objective function is unavailable or is too expensive to calculate. Under "'light-tail"' assumption for stochastic subgradient and in general case with moderate large deviation probability, we show that parallelization combined with averaging gives bounds for probability of large deviation similar to a serial method. Thus, in these cases, one can benefit from parallel computations and reduce the computational time without loss in the solution quality.

Keywords Stochastic Convex Optimization, Probability of Large Deviation, Mirror Descent, Parallel Algorithm

Mathematics Subject Classification 90C15, 90C25

1 Introduction

We consider the following general stochastic optimization problem over a convex compact set Q:

$$\min_{x \in Q \subset E} \{f(x) := \mathbb{E}_\xi [f(x, \xi)]\}, \tag{1}$$

where E is a finite-dimensional real vector space, ξ is a random vector, $f(x, \xi)$ is a closed convex function w.r.t x for a.e. ξ, and x and ξ are independent. Under these assumptions, this problem is a convex optimization problem.

Our main goal is to approximately solve this problem using some algorithm. Usually, in the stochastic optimization literature [1, 2], two measures for the quality of an approximate solution $\bar{x}$ are considered. The first is expectation of the objective residual. In this case, $\bar{x}$ is an ε-solution of (1) for $\varepsilon > 0$ iff $\mathbb{E}f(\bar{x}) - f_* \leq \varepsilon$, where f_* is the optimal value in (1) and expectation is taken with respect to all the randomness arising in the algorithmic process. The second is bound for probability of large deviation of the objective residual. In this case

[1]corresponding author, pavel.dvurechensky@wias-berlin.de, Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstr. 39, 10117 Berlin, Germany; Institute for Information Transmission Problems RAS, Bolshoy Karetny per. 19, build.1, Moscow, Russia 127051;

[2]gasnikov.av@mipt.ru, Moscow Institute of Physics and Technology, 9 Institutskiy per., Dolgoprudny, Moscow Region, Russia 141700; Institute for Information Transmission Problems RAS, Bolshoy Karetny per. 19, build.1, Moscow, Russia 127051;

[3]a.lagunovskaya@phystech.edu, Moscow Institute of Physics and Technology, 9 Institutskiy per., Dolgoprudny, Moscow Region, Russia 141700;

$\bar{x}$, is an (ε, σ)-solution of (1) for $\varepsilon > 0$, $\sigma \in (0,1)$ iff $\mathbb{P}\{f(\bar{x}) - f_* > \varepsilon\} \leq \sigma$. We mainly focus on the latter quality measure in this paper.

It is known that, if the value of the objective function is available (e.g. in randomized methods [3]), one can obtain in parallel logarithmic in σ^{-1} number of independent ε-solutions. Then, the solution with the minimum function value is an (ε, σ)-solution. Nevertheless, the value of the objective function can be unavailable or too expensive to calculate or approximate. The latter is easy to imagine since, for $\xi \in \mathbb{R}^p$, the computational effort for calculation of $\bar{f}$ s.t. $|\bar{f} - f(x)| \leq \delta$ can amount up to $O\left(\delta^{-p}\right)$ calculations of $f(x,\xi)$ at different ξ. Our goal is to propose a technique, which allows to obtain an (ε, σ)-solution based on a number computed in parallel ε-solutions without calculation of the function $f(x)$ value.

Our approach is based on Stochastic Mirror Descent algorithm [2]. It turns out that, under some mild assumptions, an ε-solution to (1), obtained by Stochastic Mirror Descent, is also an $(\tilde{\varepsilon}, \sigma)$-solution. We use this fact and calculate in parallel logarithmic in σ^{-1} number of independent ε-solutions, average them and prove that this average is an (ε, σ)-solution. Thus, we can benefit from parallelization and reduce the computation time without any loss in the solution quality.

2 Stochastic Mirror Descent

This section is devoted to description of Stochastic Mirror Descent (SMD) [2, 3] and its convergence properties in terms of expectation of the objective residual and also in terms of probability of large deviation of this residual. These convergence results provide the basis of our approach for constructing an (ε, σ)-solution by parallelization. Let us choose some norm $\|\cdot\|$ on E and denote the conjugate norm by $\|\cdot\|_*$. We assume that, at any point $x \in Q$, a stochastic subgradient $\nabla_x f(x,\xi)$ of $f(x)$ is available and satisfies

$$\mathbb{E}_\xi \left[\nabla_x f(x,\xi)\right] \in \partial f(x), \quad \mathbb{E}_\xi \left[\|\nabla_x f(x,\xi)\|_*^2\right] \leq M^2 \tag{2}$$

for some constant $M > 0$. We choose a prox-function $d(x)$ which is 1-strongly convex in $\|\cdot\|$. Let $x^0 = \arg\min_{x \in Q} d(x)$. W.l.o.g. we assume that $d(x^0) = 0$. The algorithm uses Bregman's divergence $V_z(x) = d(x) - d(z) - \langle \nabla d(z), x - z \rangle$. Let x_* be a solution of (1), R be a number s.t. $V_{x^0}(x_*) \leq R^2$, and $\bar{R}$ be a number s.t. $\max_{x \in Q} V_x(x_*) \leq \bar{R}$.

Stochastic Mirror Descent [2, 3] iterates as follows, starting from $x^0 \in Q$,

$$x^{k+1} = \mathrm{Mirr}_{x^k}\left(h \nabla_x f\left(x^k, \xi^k\right)\right), \quad \mathrm{Mirr}_{x^k}(v) := \arg\min_{x \in Q} \left\{\left\langle v, x - x^k \right\rangle + V_{x^k}(x)\right\}, \tag{3}$$

where $h > 0$ is the stepsize, $\left\{\xi^k\right\}_{k \geq 0}$ is an i.i.d. sample of ξ. The main property of SMD-step [3] is

$$2V_{x^{k+1}}(x) \leq 2V_{x^k}(x) + 2h\left\langle \nabla_x f\left(x^k, \xi^k\right), x - x^k \right\rangle + h^2 \left\|\nabla_x f\left(x^k, \xi^k\right)\right\|_*^2, \quad \forall x \in Q.$$

Further, using convexity of $f(x)$, for any $\nabla f(x^k) \in \partial f(x^k)$ and $x \in Q$,

$$f\left(x^k\right) - f(x) \leq \left\langle \nabla f(x^k), x^k - x \right\rangle \leq \left\langle \nabla f(x^k) - \nabla_x f\left(x^k, \xi^k\right), x^k - x \right\rangle +$$

$$+ \frac{1}{h}\left(V_{x^k}(x) - V_{x^{k+1}}(x)\right) + \frac{h}{2}\left\|\nabla_x f\left(x^k, \xi^k\right)\right\|_*^2.$$

Taking conditional expectation w.r.t. $\xi^1, ..., \xi^{k-1}$ and using (2), one obtains

$$f\left(x^k\right) - f\left(x\right) \leq \frac{1}{h}\left(V_{x^k}\left(x\right) - \mathbb{E}\left[V_{x^{k+1}}\left(x\right) | \, \xi^1, ..., \xi^{k-1}\right]\right) + \frac{h}{2} \underbrace{\mathbb{E}\left[\left\|\nabla_x f\left(x^k, \xi^k\right)\right\|_*^2 \middle| \, \xi^1, ..., \xi^{k-1}\right]}_{\overset{(2)}{\leq} M^2}.$$

Since $\left\{\xi^k\right\}_{k\geq 0}$ is an i.i.d. sample, taking full expectation from the both sides of these inequalities for $k = 0, ..., N-1$, summing them up and taking $x = x_*$ gives, by convexity of $f\left(x\right)$,

$$\mathbb{E}\left[f\left(\bar{x}^N\right)\right] - f_* \leq \mathbb{E}\frac{1}{N}\sum_{k=0}^{N-1} f(x^k) - f_* \leq \frac{1}{hN} V_{x^0}\left(x_*\right) + \frac{M^2 h}{2} \leq \sqrt{\frac{2M^2R^2}{N}},$$

where

$$R \quad \text{is s.t.} \quad V_{x^0}\left(x_*\right) \leq R^2, \quad \bar{x}^N := \frac{1}{N}\sum_{k=0}^{N-1} x^k, \quad h = \frac{R}{M}\sqrt{\frac{2}{N}}. \tag{4}$$

Choosing

$$N = \left\lceil \frac{2M^2R^2}{\varepsilon^2} \right\rceil, \tag{5}$$

we obtain that $\bar{x}^N$ satisfies $\mathbb{E}\left[f\left(\bar{x}^N\right)\right] - f_* \leq \varepsilon$ and, hence, is an ε-solution. Note that this bound for N is optimal [3] up to a multiplicative constant factor for the class of convex stochastic programming problems (1) with a.e. bounded stochastic subgradients.

It turns out that it is possible to prove bounds for probability of large deviation for $f\left(\bar{x}^N\right) - f_*$.

Proposition 1 ([2, 5, 6]). *Assume that one of the following assumptions holds. a)* $\left\|\nabla_x f\left(x, \xi\right)\right\|_* \leq M$ *for a.e.* ξ*; b)* $\mathbb{E}_\xi\left(\exp\left(\left\|\nabla_x f\left(x, \xi\right)\right\|_*^2 \Big/ M^2\right)\right) \leq \exp\left(1\right)$ *and* $\ln \sigma^{-1} \ll N$*; c) There exists some* $\alpha > 2$ *s.t., for all,* $t \geq 0$, $\mathbb{P}\left(\frac{\left\|\nabla f(x,\xi)\right\|_2^2}{M^2} \geq t\right) \leq \frac{1}{(t+1)^\alpha}$, *and* $\sigma^{-1/(\alpha-1)} \ll N$*. Then the point* $\bar{x}^N$ *generated by SMD* (3), (4) *satisfies*

$$\mathbb{P}\left\{f\left(\bar{x}^N\right) - f_* \leq \frac{C_1 M}{\sqrt{N}}\left(R + C_2 \bar{R}\sqrt{\ln\left(1/\sigma\right)}\right)\right\} \geq 1 - \sigma, \tag{6}$$

where in the case a) $C_1 = \sqrt{2}$, $C_2 = 2\sqrt{2}$*; in the case b)* $C_1 = C_2 = 2\sqrt{2}$*; in the case c)* $C_1 = C_1(\alpha)$, $C_2 = 1$.

Corollary 1. *Let any of three assumptions of Proposition 1 hold. Choose* $N = \left\lceil \frac{CM^2\bar{R}^2}{\varepsilon^2} \right\rceil$, *where the constant* C *depends on* C_1, C_2*. Then the point* $\bar{x}^N$ *generated by SMD* (3), (4) *satisfies, for any* $c \geq 0$,

$$\mathbb{P}\left\{f\left(\bar{x}^N\right) - f_* \geq c\right\} \leq \mathbb{P}\left\{\eta \geq c\right\}, \tag{7}$$

where $\eta \in \mathrm{N}\left(\varepsilon, \varepsilon^2\right)$ *– normal random variable with mean* ε *and variance* ε^2.

3 Parallelization and bounds for probability of large deviation

In this section, we first discuss a known way to obtain an (ε,σ)-solution using a number of ε-solutions calculated in parallel. Then, we suggest a new way of doing this without calculation of the objective $f(x)$ value, state and prove the main result.

Assume that $\bar{x}^N$ is an $\varepsilon/2$-solution, obtained by SMD (3), (4) with $N = \left\lceil \frac{8M^2R^2}{\varepsilon^2} \right\rceil$. Then, using the Markov inequality [5], we obtain

$$\mathbb{P}\left(f\left(\bar{x}^N\right) - f_* \geq \varepsilon\right) \leq \frac{\mathbb{E}\left[f\left(\bar{x}^N\right)\right] - f_*}{\varepsilon} \leq \frac{1}{2}.$$

If one calculates in parallel $K = \left\lceil \log_2\left(\sigma^{-1}\right) \right\rceil$ independent SMD $\varepsilon/2$-solutions $\left\{\bar{x}^{N,i}\right\}_{i=1}^K$ and chooses the one $\bar{x}^N_{\min}$ which minimizes $f\left(\bar{x}^{N,i}\right)$, then, in total $\left\lceil \frac{8M^2R^2}{\varepsilon^2} \right\rceil \left\lceil \log_2\left(\sigma^{-1}\right) \right\rceil$ calculations of stochastic subgradient, one obtains $\mathbb{P}\left(f\left(\bar{x}^N_{\min}\right) - f_* \geq \varepsilon\right) \leq \sigma$. Thus, $\bar{x}^N_{\min}$ is an (ε,σ)-solution of (1). The crucial point here is the possibility to calculate the value of the function $f(x)$.

We now suggest a technique which does not rely on the assumption of the function $f(x)$ value availability. This assumption may not hold [1] in many real stochastic programming problems, e.g. in maximum likelihood approach used in mathematical statistics.

Theorem 1. *Let any of three assumptions of Proposition 1 hold. Let* $K = \left\lceil 2\ln\left(\sigma^{-1}\right) \right\rceil$ *and* $\left\{\bar{x}^{N,i}\right\}_{i=1}^K$ *be independent points obtained by SMD* (3), (4) *with* $N = \left\lceil \frac{4CM^2\bar{R}^2}{\varepsilon^2} \right\rceil$. *Then the point* $\bar{x}^K = \frac{1}{K}\sum\limits_{i=1}^K \bar{x}^{N,i}$ *is an* (ε,σ)*-solution of* (1).

Proof Using Corollary 1 with $c = \varepsilon$, we obtain, for all $i = 1, ..., K$,

$$\mathbb{P}\left\{f\left(\bar{x}^{N,i}\right) - f_* \geq \varepsilon\right\} \leq \mathbb{P}\left\{\eta_i \geq \varepsilon\right\}, \quad \eta_i \in \mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4}\right). \tag{8}$$

By convexity of $f(x)$, since $\left\{\bar{x}^{N,i}\right\}_{i=1}^K$ are i.i.d., $\{\eta_i\}_{i=1}^K$ are i.i.d.,

$$\mathbb{P}\left\{f\left(\bar{x}^K\right) - f_* \geq \varepsilon\right\} \leq \mathbb{P}\left\{\frac{1}{K}\sum_{k=i}^K \left(f\left(\bar{x}^{N,i}\right) - f_*\right) \geq \varepsilon\right\} \overset{(8)}{\leq} \tag{9}$$

$$\mathbb{P}\left\{\frac{1}{K}\sum_{i=1}^K \eta_i \geq \frac{\varepsilon}{2} + \frac{\varepsilon}{2}\right\} = \mathbb{P}\left\{\left(\frac{1}{K}\sum_{i=1}^K \eta_i\right) - \frac{\varepsilon}{2} \geq \frac{\varepsilon}{2}\right\} = \mathbb{P}\left\{\eta - \frac{\varepsilon}{2} \geq \frac{\varepsilon}{2}\sqrt{K}\right\} \leq \sigma, \tag{10}$$

where $\eta \in \mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4}\right)$. Here we used well-known facts about the properties of sum of independent normal random variables [4]

$$\mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4}\right) + ... + \mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4}\right) \overset{d}{=} \mathrm{N}\left(\frac{K\varepsilon}{2}, \frac{K\varepsilon^2}{4}\right), \quad \frac{1}{K}\mathrm{N}\left(\frac{K\varepsilon}{2}, \frac{K\varepsilon^2}{4}\right) \overset{d}{=} \mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4K}\right),$$

that $K = \left\lceil 2\ln\left(\sigma^{-1}\right) \right\rceil$ and, for $\eta \in \mathrm{N}\left(\frac{\varepsilon}{2}, \frac{\varepsilon^2}{4}\right)$, $\mathbb{P}\left\{\eta - \frac{\varepsilon}{2} \geq \frac{\varepsilon}{2}\sqrt{2\ln\left(\sigma^{-1}\right)}\right\} \leq \sigma$. Thus, the point $\bar{x}^K$ is an (ε,σ)-solution of (1). □

To sum up, under some mild assumptions, but without calculation of the objective function value, we propose a way to obtain an (ε, σ)-solution using a number of ε-solutions calculated in parallel. This approach allows to reduce the computational time without any loss in the solution quality. At the same time, we answer to the question of Yu. Nesterov [7]. The question can be stated as follows. When the quality of a solution of problem (1) obtained by one wise old man, thinking for $\Theta\left(M^2\bar{R}^2 \ln\left(\sigma^{-1}\right)/\varepsilon^2\right)$ days, is the same as obtained by $\Theta\left(\ln\left(\sigma^{-1}\right)\right)$ experts, each thinking for $\Theta\left(M^2\bar{R}^2/\varepsilon^2\right)$ days? Our answer is that the quality is equivalent under any of three assumptions of Proposition 1.

References

[1] *Shapiro A., Dentcheva D., Ruszczynski A.* Lecture on stochastic programming. Modeling and theory. – MPS-SIAM series on Optimization, 2014.

[2] *Nemirovski A., Juditsky A., Lan, G., Shapiro, A.* Robust stochastic approximation approach to stochastic programming. SIAM J. Optim. 2009, V. 19. P.1574-1609.

[3] *Nemirovski A.* Lectures on modern convex optimization analysis, algorithms, and engineering applications. Philadelphia: SIAM, 2013.

[4] *Durrett R.* Probability and examples. – Cambridge University Press, 2010.

[5] *Guiges V., Juditsky A., Nemirovski A.* Non-asymptotic confidence bounds for the optimal value of a stochastic program // e-print, 2016 arXiv:1601.07592

[6] *Gasnikov A.V.* Searching equilibriums in large transport networks. Doctoral thesis. – MIPT, 2016. [in Russian] arXiv:1607.03142

[7] *Nesterov Yu., Vial J.-Ph.* Confidence level solution for stochastic programming // Automatica. 2008. V. 44. no. 6. P. 1559–1568.